A locally-adaptive metric for contrast-fusion of noisy multimodal imagery

Diego A. Socolinsky
Equinox Corporation
9 West 57th Street, New York, NY, USA
diego@equinoxsensors.com
phone: 212-421-2999 fax: 212-421-7004

Abstract—

We propose a new technique for reducing the effect of noise on contrast fusion of multimodal imagery, yielding higher quality results than can be obtained with previous methods. We rely on local non-parametric estimation of band gradient entropies as a relative measure of geometric structure. This work builds upon previous work by the author, and is exemplified with remote sensing applications. Quantitative measures of performance are given on ground-truth data, which indicate the advantages of the new technique over existing approaches.

Keywords— **Multispectral image processing, image fusion, remote sensing, adaptive noise estimation, entropy estimation.**

I. Introduction

Increasing availability of co-registered multimodal imagery from diverse sources such as magnetic resonance scanners, aerial and earth orbiting sensors has spurred development of numerous techniques for image fusion and visualization [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12]. These methods can be roughly categorized in two classes, those which work on the zero order properties of the images, and those that use higher order information such as first or second derivatives. We will refer to the later as contrast fusion methods, since they attempt to reproduce a combination of local contrast from each modality in the fused image product. These methods appear to hold the greatest promise, despite their computational burden, since they are well-adapted to the physiological basis of contrast perception in the low-level human visual system [13], [14].

Contrast-based fusion techniques typically rely on the absolute magnitude of derivatives as a means of determining which bands of a multimodal image dominate the result at a given point. A common rule is to choose the maximum contrast among the bands and prescribe that as the contrast for the fused image, with the rationale that large contrast correlates with visually relevant image features. While this assumption may hold for noise-free images, it becomes false when one or several bands are corrupted by noise. In that situation, large contrast may correspond to variation due to noise, and not to underlying image features. Thus any fusion rule that assumes all contrast to be due to valid signal variation will inevitably produce unsatisfactory results in the presence of noise. The present article proposes a new technique for minimizing the adverse effects of noise on contrast fusion algorithms, thereby rendering them more widely applicable. Most or all previous contrast fusion algorithms could be adapted to take advantage of this new technique. However, the algorithms previously developed by the author and his co-workers are especially well-suited to incorporate the ideas introduced in this paper, and therefore these algorithms are used as a model.

This research was partially funded by DARPA contract #DAAH01-99-C-0047, and AFOSR contract #F49620-98-C-0047. This article was completed while the author was a visiting scholar at the Department of Mathematical Sciences, School of Engineering, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218.

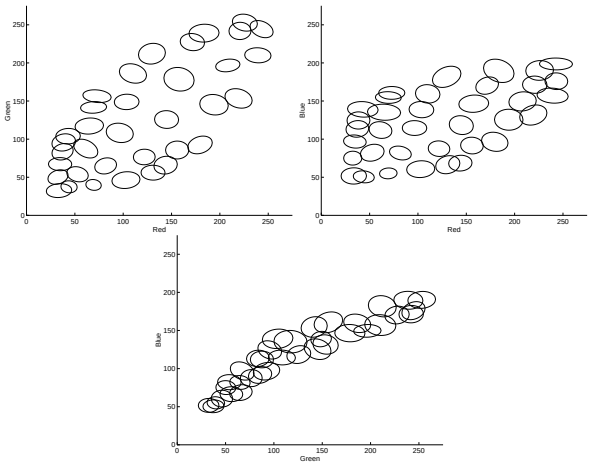

Fig. 1. Planar projections of the sensor noise equal probability ellipses, computed for a Cohu 2200 color camera, in RGB coordinates.

The outline of this article is as follows. In section II we briefly review the contrast fusion algorithm proposed in [5], [13]. Section III describes a technique for including information about the sensor noise characteristics into the fusion process. A locally adaptive merging rule that takes into account absolute contrast magnitude and image noise is introduced in section IV. Section V shows experimental results obtained on various images. Finally section VI provides some conclusions and outlines directions for future improvement.

II. Contrast fusion model

The method introduced below applies to several contrast fusion models, such as those based on Laplacian pyramids and wavelet transforms. For simplicity we introduce it in the context of the model in [5] and the extensions in [13], which we review in this section.

Let $\Omega \subset \mathbb{R}^2$ denote the image domain (usually a rectangle), and consider the trivial Riemannian vector bundle $P = (\Omega, \Omega \times \mathbb{R}^n, g)$, where Ω is the base space, $\Omega \times \mathbb{R}^n$ is the total space, g is a fiber-wise constant metric, and the projection map is simply $\pi(x, y, v) = (x, y)$, for $(x, y) \in \Omega$, $v \in \mathbb{R}^n$. Now, define an n-band image as a section of V, that is a map of the form $\Omega \ni (x, y) \overset{S}{\mapsto} (x, y, s(x, y))$, with $s : \Omega \to \mathbb{R}^n$. One may then write $S = Id \otimes s$, where Id denotes the identity map on $\mathbb{R}^2$. The corresponding *contrast form* is defined as $\chi = S^* h - \delta$, where δ is the Euclidean metric on $\mathbb{R}^2$, and S^* denotes the pullback map induced on differential forms on $\mathbb{R}^n$ by S. This is a bilinear form, which is usually degenerate since s may have vanishing differential.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
2000		**00-00-2000 to 00-00-2000**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
A locally-adaptive metric for contrast-fusion of noisy multimodal imagery	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Equinox Corporation,9 West 57th Street,New York,NY,10019**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**		**6**	

In coordinates, the contrast form is given by

$$\chi_{ij}(p) = \sum_{k,l=1}^{n} g_{kl}(S(p)) \frac{\partial s^k(p)}{\partial x_i} \frac{\partial s^l(p)}{\partial x_j} - \delta_{ij},$$

for $1 \leq i, j \leq 2$ and $p \in \Omega$.

The *contrast vector field* $V \in T\Omega$ is defined so that its magnitude and direction correspond with the largest eigenvalue of χ and its eigenspace, and so that its orientation agrees with that of the gradient of an auxiliary function $\phi : \Omega \to \mathbb{R}$, usually chosen to be $\phi = \sum_{i=1}^{n} s^k$ (but see [13] for other choices). It is straightforward to verify that if $n = 1$ and the metric on the bundle P is globally Euclidean, then $V = \nabla s$, hence the contrast vector field reduces to the standard notion of first order contrast for grayscale images. In the case $n > 1$, the contrast vector field encodes the combination of contrast in all bands of the image S. Therefore, a grayscale image, given as an intensity map $f : \Omega \to \mathbb{R}$, whose gradient equals V is a fused realization of the contrast in all bands of S. Note that the bundle metric g controls how contrast from each band contributes to the contrast of the fused image, and furthermore how contrast interacts among bands. If g is chosen to be Euclidean on each fiber, then all bands contribute equally and do not interact with each other non-linearly.

It is normally not possible to solve the equation $\nabla f = V$ as a means of constructing the fused image, so we resort to finding the best L^2-approximate solution, which is characterized (up to a constant) as a solution of the Neumann problem

$$\begin{cases} \Delta f = \operatorname{div} V, & \text{on } \Omega, \\ \nabla f \cdot \vec{n} = 0, & \text{on } \partial\Omega. \end{cases}$$

This is the desired fusion of the multiple bands in the original image. For connections with wavelet-based fusion algorithms and L^1-approximate solutions see [13]. There are many efficient numerical methods for solving this equation, and we will not discuss this topic in the current article (see [5] and the references therein).

III. Sensor dependent metric

Wolff and Socolinsky [15] proposed a method for handling sensor noise in the context of edge detection in multispectral images, which can also be used to account for sensor characteristics within the fusion model of section II. Their approach is based on a generalization of the MacAdams color discrimination metric of the human eye [16], to artificial camera sensors, and models the inherent noise properties of the sensor as a function of position in photometric space. Let us explain the idea for a standard 3-color camera. Let $X = (x_1, x_2, x_3)$ be the color coordinates, in some fixed coordinate system, of an incoming light stimulus. The camera response $C(X)$ for a stimulus of color X is a random variable, which they assume to be Gaussianly distributed with mean X and covariance $\Sigma(X)$. The parameter $\Sigma(X)$ measures the reliability of the sensor response for stimuli with color coordinates X. Thus, it is natural to define a global metric on photometric space, independent of spatial coordinates, by

$$g(X) = \Sigma^{-1}(X).$$

Using this metric, contrast from each band is weighted according to the noise characteristics of the sensor for each particular color, whereby noisier bands will contribute less to the contrast form than bands with lower variance, in which we have higher confidence. Figure 1 shows planar projections of the equi-probability ellipses computed for a Cohu 2200 color camera, represented in RGB coordinates. The main problem with this approach is that in order to compute the metric we must have access to the sensor and a rather substantial experiment must be carried out. Even then, noise properties of the sensor may be dependent on many changing factors, such as environmental conditions, and these should be taken into account. Consequently, this construction may be of theoretical, more so than computational, interest.

IV. Locally adaptive metric

Not all image noise can be accounted for by the sensor dependent metric from the previous section. Furthermore, we often have no access to the sensor or cannot account for all variables which figure into the noise metric. Under such circumstances, it may be advantageous to choose the metric adaptively, so as to maximize some information criterion. It is not possible to separate signal from noise within a single band, but having multiple bands we may hope to determine to what degree contrast from one band is more 'structured' than that from another, and then use this knowledge to decide on an appropriate weighting.

It is tempting to consider the local mutual information between the gradients of different bands as a measure of consistency. However, it is easy to see that mutual information does not give us the kind of measure we are looking for. In fact, recall that for independent random variables, the mutual information is zero, but the distribution of gradients in different bands may be uncorrelated for different reasons. Consider for example a three-band image set consisting of a thermal-infrared band, a range scanner band and an intensified near-infrared band, acquired under extremely low light conditions. In this situation, the intensified band will be very noisy, while the other two bands will not be adversely affected by the illumination. Now certain pixels will show strong edges in the range band, due to depth discontinuities without temperature change, and at other pixels the reverse will be true due to thermal variations with no change in depth. For a neighborhood of these pixels the mutual information between the respective gradients may be low, but this is exactly the kind of situation where we want to take into account the input from both bands (or at least the one with strongest gradients). On the other hand, at certain image pixels we will have strong gradients in the intensified near-IR band due entirely to noise and weaker but relevant edges in another band due to the true signal. Here the mutual information will still be low, but both bands should not contribute equally to the mixture. Furthermore, the band with largest gradients should not be weighted more heavily as it is noisier.

More generally, one cannot appeal to methods which exploit any sort of consistency or correlation across bands, since in the most interesting image fusion scenarios that correlation may be null. Indeed, one of the most profitable applications of image fusion are those in which different image bands carry complementary and uncorrelated information. In these cases several bands are **necessary** to get a complete informative picture. What we would like to do is combine the band gradients so that the image obtained from the resulting contrast form takes optimal advantage of each band, emphasizing structured bands and de-emphasizing noisy ones. We now

propose a method which yields good results, but has no provably optimal properties. The optimal choice of metric, if such an object can in fact be characterized, is an open problem.

Entropy-based techniques have been used in image processing for tasks such as reconstruction from incomplete or noisy data [17] and removal of motion artifacts in MRI [18]. Let $S = Id \otimes s$, with $s : \Omega \to \mathbb{R}^n$ be an n-band image. For each band s^k, $k = 1 \dots n$, of s and each pixel $(x, y) \in \Omega$, we wish to consider the entropy of the random vector ∇s^k in a discrete pixel neighborhood of (x, y).

For simplicity, consider a square pixel neighborhood of odd side length w , centered at the pixel $(x, y) \in \Omega$, and denote it by $U_w(x, y)$. Circular neighborhoods can also be used with similar results and more computational effort. We estimate the local probability density of ∇s^k over $U_w(x, y)$ nonparametrically by

$$\hat{Q}(v) = \frac{1}{w^2\,\lambda} \sum_{p \in U_w(x,y)} K(\frac{|v - \nabla^k(p)|}{\lambda}), \quad v \in \mathbb{R}^2,$$

which gives us an estimate for the local entropy

$$H_w^k(x, y) = -\int_{\mathbb{R}^2} \hat{Q}\ \log \hat{Q},$$

where, as is customary, the logarithm above is taken to be zero if $\hat{Q}$ vanishes, K is a smoothing kernel (Gaussian in our numerical implementation), and λ is a fixed bandwidth. It follows from the definition of entropy that $0 \leq H_w^k(x, y) \leq 2 \log w$. From this computation, we obtain n maps

$$\tilde{H}_w^k : \Omega \to [0, w^2], \qquad k = 1 \dots n$$

defined by

$$\tilde{H}_w^k(x, y) = e^{2 \log w - H_w^k(x,y)}.$$

Now, we define a spatially varying, locally adaptive, diagonal metric on the trivial bundle $\Omega \times \mathbb{R}^n$ by

$$g_{ij}(x, y) = \begin{cases} \frac{n\tilde{H}_w^i(x,y)}{\sum_{k=1}^n \tilde{H}_w^k(x,y)} & \text{if } i = j, \\ 0 & \text{otherwise} \end{cases} \qquad (1)$$

It behooves us now to provide some intuition for the nature of the entropy-weighted metric (1). Consider the difference between an image neighborhood with geometric structure and one which is dominated by noise. In the structured case, we expect the gradients to be concentrated in a few directions, corresponding to edge orientations, and magnitudes, corresponding to edge points versus non-edge points. For a noisy neighborhood, we should expect the gradients to be distributed randomly, spread over a large range of magnitudes and orientations. Therefore, the entropy of the gradient, as a two-dimensional random variable, in an image neighborhood can be expected to be lower for well-structured neighborhoods than for ones dominated by noise. Note that we cannot use the entropy of the raw grayvalues from each band as a weighting criterion. In fact, the local entropy of the grayvalues is invariant under spatial rearrangement of the pixels in the estimation neighborhood, and thus does not convey information about the geometric structure of the image over that neighborhood. On the other hand, local gradient entropy is not invariant under rearrangement, since the spatial location of pixels within the neighborhood is used to compute the gradients, and thus different pixel arrangements yield different gradient distributions.

Figure 2 shows an example of the distribution of gradients for a noisy neighborhood versus a structured one, as pictured in figure 2 (left). One can easily see that the distribution of gradients for the unstructured neighborhood is rather uniform. The geometric structure of the second example neighborhood is apparent from the tight clustering of gradient vectors into two groups, those corresponding to the vertical edge through the middle of the neighborhood, and the rest. The metric in (1) exploits this fact by favoring bands with lower local entropy. If all bands have approximately the same gradient entropy in a neighborhood of a given pixel, then no decision can be made and the metric defaults to Euclidean at that point. The free parameters in the metric are the size of the neighborhood over which we estimate the local entropy, and the bandwidth of the kernel estimator, both of which are scale parameters. We use a fixed neighborhood width at every pixel of every band. It is very possible that an adaptive width entropy estimation could yield better results, but we have been unable to find an effective adaptation rule, since that rule should somehow depend on the noise properties of the image bands in that neighborhood and that is precisely what we are trying to estimate in the first place.

The entropy metric in (1) provides a pointwise measure of the confidence we have in the contrast from a given band at that point being due to a structured signal and not random noise. Since only the information in the image being processed is used, we can only provide a relative measure, where the confidence for one band is relative to the properties of the other bands in the same neighborhood. This construction effectively implements Ockham's razor, favoring the simplest alternative. Note that this is independent of the contrast present in each band at a given pixel. Individual band contrasts still compound to create the contrast vector field via the contrast form as in section II, but their relative weighting is given by our entropy metric.

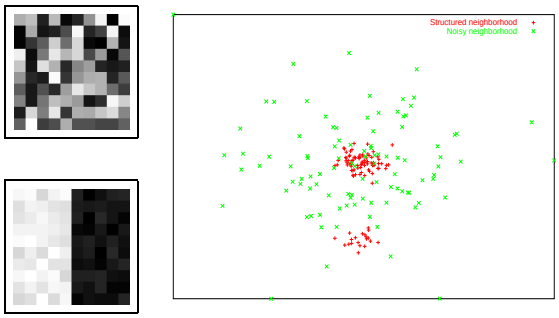

Fig. 2. Example gradient distributions for noisy (green) and structured (red) neighborhoods at left.

V. Experimental Results

Let us first look at an example from remote sensing. Figure 3 shows three bands of a multispectral aerial image acquired over Hawaii with dynamic range $[0, 255]$, and a fourth artificial band composed of Gaussian distributed noise with mean 127 and standard deviation of 40. If we fuse these four bands with the method in section II using a Euclidean metric, we obtain the result in Figure 4(a). Note how all the noise in the artificial band has been transferred to the fused image, as it

Fig. 3. Three bands of a multispectral aerial image acquired over Hawaii (courtesy of Space Imaging Corporation), plus an artificial band of Gaussian noise.

was equally weighted with respect to the other three bands. Any other contrast fusion method that does not explicitly take into account noise will produce comparable results. If instead of a Euclidean metric, we use the one constructed in (1), then we obtain the fused image in Figure 4(b), in which the noisy band has been automatically de-emphasized. A side-by-side comparison of subregions of the images in 4(a) and 4(b) is provided in Figure 5. It is evident that much less noise has been transferred to the fused image through the use of the entropy-weighted metric. In fact, much detail is destroyed by the effect of noise in the Euclidean fusion, especially in areas with weak edges in the clean bands which nonetheless represent true structure. This is especially visible in the region obscured by cloud shadows magnified at the bottom of figure 5. The fusion with entropy-weighted metric successfully ignores contrast due to noise in the artificial band, when more structured contrast is present in the other bands. From a quantitative point of view, the mean signal-to-noise ratio for the Euclidean fusion is 40.47dB and for the entropy-weighted fusion it is 58.12dB, an improvement of over 43%.

To gain some more intuition, we can look at the components of the entropy-weighted metric for the example image in figure 3. Figure 6 shows a plot of the diagonal metric coefficients along a horizontal scan-line three-quarters of the way down the image in figure 3, along with a comparison of the average coefficient of the first through third bands versus the coefficient of the fourth (noise) band. We see that whenever the first three bands show more structure at the scale we use for estimating the local entropy, the are favored heavily over the noisy band. On the right side of the plot, which corresponds to the heterogeneous region of tightly packed houses on the lower right side of the image, the weight on the noise band becomes comparable to the other three weights, as the local entropy estimation cannot differentiate structure from noise at the given scale.

A fair measure of performance can be obtained from a simple artificial experiment, similar to the example above. We took a noise-free image (the standard Lena test image), created a second band of random noise with the same mean and standard deviation. Then, we computed the fusion of these two bands using the algorithm in section II, both with a Euclidean metric and the locally-adaptive metric in section IV, as well as with the wavelet method in [8], using a Haar basis. The best possible result is obviously to recover the original noise-free image, with no corruption from the artificial noise band. Measuring the departure from the optimal result with the mean squared signal-to-noise ratio, we obtain the results in table I, corresponding to the images in figure 7. One can easily see that the local-entropy-based metric is superior to the Euclidean one, and both outperform the (Euclidean) wavelet

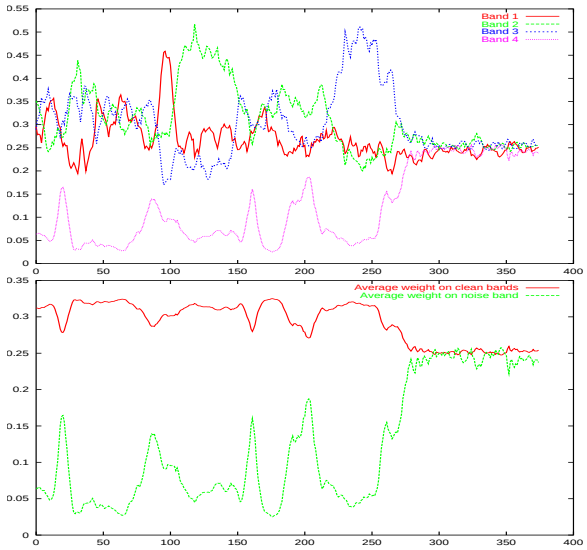

Fig. 6. Top: Diagonal components of the entropy-weighted metric along a horizontal scan-line of the image in figure 3. Bottom: Average metric coefficient on clean bands versus coefficient on noise band.

fusion in this example. The difference between the Euclidean fusion from section II and the wavelet result should not necessarily be interpreted as indication that one method is superior to the other for general 2-band images. In this example, the mean value of the metric component corresponding to the Lena image is 0.83, while the mean metric component on the noise band is 0.17. Not surprisingly, referring to table I we see that $0.83/0.17 \approx 97.47/20.10$, which indicates that the capacity of the metric to distinguish structure from noise is almost directly transferred to the fusion algorithm.

	Euclidean	Local Entropy
Contrast Form	20.10	97.47
Haar Wavelet	9.40	N/A

TABLE I

MEAN SQUARED SIGNAL-TO-NOISE RATIOS FOR RECONSTRUCTION OF LENA IMAGE WITH DIFFERENT METHODS AND METRICS.

In order to test the sensitivity of the local entropy estimation to different degrees of noise, we conducted the following experiment. The standard Lena image (upper left on figure 7) was corrupted with additive Gaussian noise of different standard deviations, on a grid pattern over the image and again on the complement of that pattern. Thus we obtained a synthetic 2-band image, each band of which has added noise exactly where the other does not, for each choice of standard devia-

(a) (b)

Fig. 4. Fusion of four bands in Figure 3. (a) Using a Euclidean metric. (b) Using an entropy-weighted diagonal metric with $w = 9$.

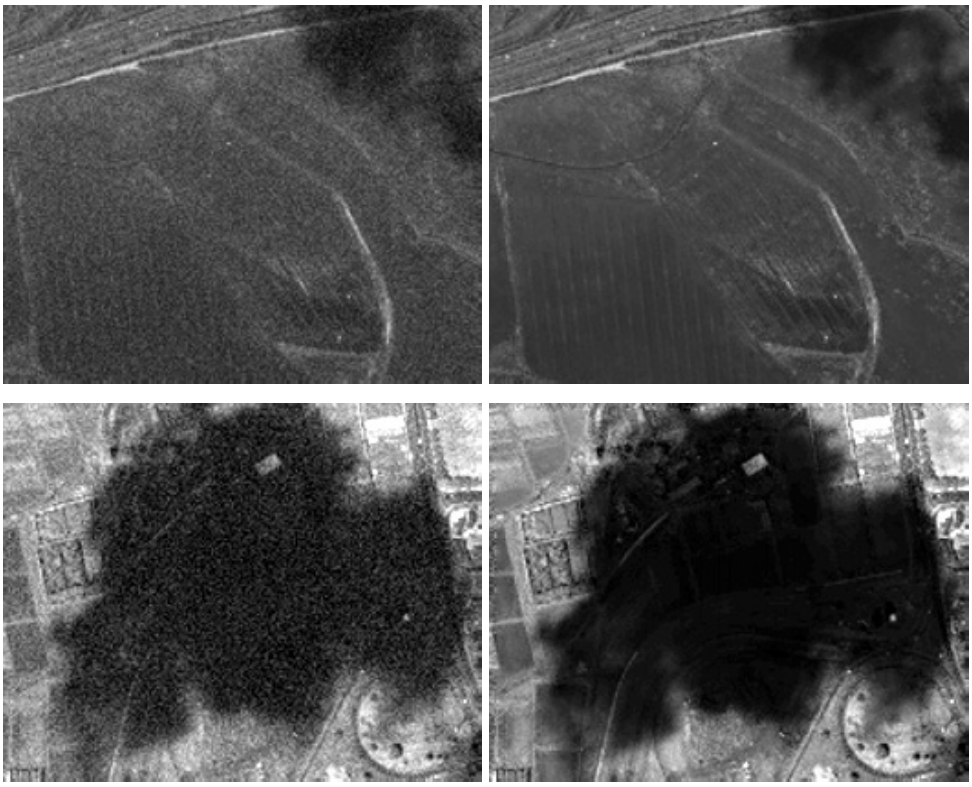

Fig. 5. Close-ups from the images in Figure 4 showing the effect of the entropy-weighted metric (bottom images are gamma-corrected to show structure in dark areas).

tion. Figure 9(a) and 9(b) show an example of this for noise of standard deviation 45 grayvalues, with the added noise highlighted in red. For each such pair we computed the metric constructed in section IV, and made a pixel-wise decision as to which band was noisier, based on the magnitude of each metric component. The proportion of correctly classified pixels as a function of the standard deviation of the added noise is plotted in figure 8 (left). Note that even at a standard deviation of 2.5 grayvalues the percentage of correctly classified pixels is approximately 89%. The majority of misclassified pixels lie on the hair, where the correct image structure is almost impossible to detect locally without semantic clues, and along the grid boundaries.

VI. Conclusion

We introduced a locally adaptive metric associated to a multimodal image, which models the relative likelihood that features from a given band are due to underlying image features, versus noise. The underlying assumption is that neighborhoods with simpler geometric structure are more likely to arise from true features, while geometrically disorganized ones are likely to be due to noise. This metric can then be instantiated into the image fusion formalism previously proposed by the author, to yield a contrast fusion model which is more robust to the presence of image noise. We presented experimental results that clearly show the effect of the adaptive metric on real-life imagery, with synthetic noise. We can numerically compute the effect of the new metric on the resulting fusion,

Fig. 7. From top left to bottom right: original image, result of Euclidean fusion, result of fusion with local-entropy-based metric, result of Haar wavelet fusion.

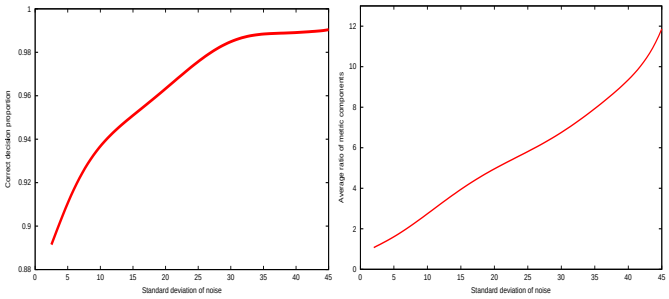

Fig. 8. Left: Proportion of correctly classified pixels as a function of the standard deviation of additive noise. Right: Ratio of metric components a a function of standard deviation of additive noise

which is seen to be very positive.

Two topics for future research should be clear from the construction above. Firstly, the metric we construct is pointwise diagonal, which means that contrasts from different bands are not allowed to interact with each other in a nonlinear fashion. The reason for placing such a restriction on the metric was simply that in the diagonal case the use of local gradient entropies becomes intuitive. However, there is no reason to believe that the ideal metric, if such an object exists, should be diagonal, and thus the construction of non-diagonal metrics that extend the results above should be explored. Secondly, as we noted in section IV, our entropy estimation relied on neighborhoods of fixed size throughout the image plane. It makes sense to think that if we knew *a priori* that a given pixel is surrounded by a geometrically structured region, then a smaller neighborhood would suffice for estimating the entropy, while larger neighborhoods would be required for pixels in noisy regions. This would allow the metric to more closely follow the local characteristics of the multiband image. However, when we have no *a priori* knowledge of the image noise, it is not clear how to vary the local neighborhood size.

Lastly, we should note that by its very nature this method is not well-adapted for processing imagery with structured or periodic noise at the estimation scale. In that case, the entropy based metric will detect the structure of the noise and mistakenly increase the weight assigned to bands containing such noise. In those situations, pre-processing each band separately with a filter in the frequency domain may yield better results.

Fig. 9. (a) and (b): Two bands of an artificial image with added Gaussian noise of standard deviation 45 grayvalues highlighted in red. (c) Misclassified pixels for image pair with noise of standard deviation 5 grayvalues.

REFERENCES

[1] A. Manduca, "Multi-spectral medical image visualization with self organizing maps," in *Proceedings IEEE International Conference on Image Processing*, 1994, pp. 633–637.

[2] L. van Ruyven A. Toet and J. Valeton, "Merging thermal and visual images by a contrast pyramid," *Optical Engineering*, vol. 28, no. 7, pp. 789–792, 1989.

[3] A. Toet, "New false color mapping for image fusion," *Optical Engineering*, vol. 35, no. 3, pp. 650–658, 1996.

[4] D. Fay et al. A. Waxman, A. Gove, "Color night vision: Opponent processing in the fusion of visible and IR imagery," *Neural Networks*, vol. 10, no. 1, pp. 1–6, 1997.

[5] D. A. Socolinsky and L. B. Wolff, "A new paradigm for multispectral image visualization and data fusion," in *Proceedings: CVPR '99*, Fort Collins, June 1999.

[6] J. Schuler M. Satyshur D. Scribner, P. Warren and M. Kruer, "Infrared color vision: An approach to sensor fusion," *Optics and Photonics News*, August 1998.

[7] G. Harikumar and Y. Bresler, "Feature extraction for exploratory visualization of vector valued imagery," *IEEE Trans. on Image Processing*, vol. 5, no. 9, pp. 1324–1334, 1996.

[8] B. S. Manjunath H. Lui and S. K. Mitra, "Multi-sensor image fusion using the wavelet transform," in *Proceedings IEEE International Conference on Image Processing*, 1994, pp. 51–55.

[9] M. Ehlers, "Multisensor image fusion techniques in remote sensing," *ISPRS J. Photogrammetry and Remote Sensing*, vol. 46, pp. 19–30, 1991.

[10] P. Burt and R. Lolczynski, "Enhanced image capture through fusion," in *Proceedings of IEEE 4th International Conference on Computer Vision*, 1993, vol. 4, pp. 173–182.

[11] S. K. Rogers T. A. Wilson and M. Kabrisky, "Perceptual-based image fusion for hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 4, pp. 1007–1017, July 1997.

[12] U. Schmiedl, D. A. Orthendahl, A. S. Mark, I. Berry, and L. Kaufman, "The utility of principal component analysis for the image display of brain lesions: A preliminary, comparative study," *Magnetic Resonance in Medicine*, vol. 4, pp. 471–486, 1987.

[13] D. A. Socolinsky, *A Variational Approach to Image Fusion*, Ph.D. thesis, The Johns Hopkins University, Baltimore, MD, 2000.

[14] R. W. Rodiek, *The vertebrate retina: principles of structure and function*, W. H. Freeman and Company, San Francisco, 1973.

[15] L. Wolff and D. Socolinsky, "Theory and analysis of color discrimination for the computation of color edges using camera sensors in machine vision," in *Proceedings: 8th Congress of the International Colour Association*, Kyoto, May 1997, pp. 515–518.

[16] G. Wyszecki and W. S. Styles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, John Wiley & Sons, New York, 1982.

[17] S. F. Gull and G. J. Daniell, "Image reconstruction from incomplete and noisy data," *Nature*, vol. 272, pp. 686–690, 1978.

[18] P. Stoyle *et al* D. Atkinson, D. Hill, "Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion," *IEEE Transactions on Medical Imaging*, vol. 16, no. 6, pp. 903–910, December 1997.

[19] D. A. Socolinsky and L. B. Wolff, "Image fusion for enhanced visualization of brain imaging," in *Proceedings: SPIE Medical Imaging '99*, San Diego, February 1999.